

The Performance of the NAS HSPs in 1st Half of 1994

Robert J. Bergeron

Report NAS-95-008 February 1995

NAS Systems Development Branch
NAS Systems Division
NASA Ames Research Center
Mail Stop 258-6
Moffett Field, CA 94035-1000

Abstract

During the first 6 months of 1994, the NAS 16-CPU Y-MP C90 Von Neumann (VN) delivered an average throughput of 4.045 GFLOPS while the ACSF 8-CPU Y-MP C90 Eagle averaged 1.658 GFLOPS. The VN rate represents a machine efficiency of 26.3% whereas the Eagle rate corresponds to a machine efficiency of 21.6%. VN displayed a greater efficiency than Eagle primarily because the stronger workload demand for its CPU cycles allowed it to devote more time to user programs and less time to idle. An additional factor increasing VN efficiency was the ability of the UNICOS 8.0 Operating System to deliver a larger fraction of CPU time to user programs. Although measurements indicate increasing vector length for both workloads, insufficient vector lengths continue to hinder HSP performance. To improve HSP performance, NAS should continue to encourage the HSP users to modify their codes to increase program vector length.

1.0 Introduction

The introduction of the C90 in March 1993 motivated the daily monitoring of the hardware performance of the NAS High Speed Processors (HSPs). The C90 Hardware Performance Monitor (HPM) continuously delivers a full 32-counter record for the workload [1]. NAS records the daily average values of all HPM counters and this paper, covering the 1st half of 1994, is the fourth in a series of reports using these counters to evaluate the performance of the NAS workload.

A NASA Ames administrative action brought the Aeronautics Consolidated Supercomputer Facility (ACSF) C90 Eagle under control of NAS in the second half of 1993. Daily HPM monitoring of the ACSF C90 Eagle began in late March of 1994. The following table presents the characteristics of the two machines.

Table 1: Characteristics of NAS HSPs

Characteristic	Unit	Von Neumann	Eagle
Serial Number	---	4012	4015
Number of CPUs	---	16	8
Clock Cycle	ns	4.167	4.167
Memory	MW	1024	128
Memory Banks	---	1024	512
Memory Bank Cycle Time	Clock Period	23	23
SSD Size	MW	1024	512
VHISPs	---	4	2
IOS	---	Model E	Model E
UNICOS	Version	8.0.1	7.C.2

This report provides tables of counter values representing the average, maximum, and minimum values from the daily reports in the first half of 1994. The NAS C90 Von Neumann(VN) provided 180 such daily reports

whereas the ACSF Eagle provided 105 daily reports since monitoring began in the middle of the first half of 1994.

The counter value tables provide performance rate data per CPU for the actual time the CPU spent executing the user programs. System throughput is derived from the CPU Floating Point Operation (FLOP) rate, the wall clock time, and the total number of CPUs. A complete explanation of all counter data occurs in [2].

To provide a feel for the daily variation in each of the counters, the report also provides the standard deviation (STD) and coefficient of variation (COV). The coefficient of variation is the ratio of the standard deviation of a quantity divided by its average value.

The report divides the 32 C90 counters into 4 functional groups: global counters, instruction holds, instruction issues, and vector operations. Sections 2 through 5 describe the results obtained from each of the four groups and compares the measurements from the two workloads.

2.0 Global Counter Data

Table 2 provides counter data giving a total counts for instructions, operations and references. The unit "M/sec" denotes "Million per sec" and the unit "avg/ref" denotes "average (conflict) per reference". The term "reference" denotes a single Cray word (8-byte) data transfer.

**Table 2: NAS C90 VN 1H94 Daily Average HPM Measurements-
Global Counters**

Measurement	Unit	Avg	STD	COV	Min	Max
CPU time	Sec	1096107.	277109.	0.253	155327.	1531277.
Instruction Issue	M/sec	53.583	3.379	0.063	43.297	61.972
Average clock periods/inst	-----	4.497	0.287	0.064	3.872	5.543
CP holding issue	Percent	69.716	2.019	0.029	63.776	75.239
Instruction buffer fetches	M/sec	0.273	0.039	0.141	0.178	0.385
Floating Pt. Ops per CPU	M/sec	267.347	22.475	0.084	218.794	346.303
Vector Fl. Pt. Ops per CPU	M/sec	263.937	22.838	0.087	213.743	344.396
CPU memory references	M/sec	267.175	18.371	0.069	217.479	330.843
CPU memory conflicts	Avg/ref	0.300	0.049	0.164	0.226	0.643
VEC memory references	M/sec	262.550	18.936	0.072	210.648	327.693
B/T memory references	M/sec	1.262	0.219	0.173	0.712	1.974
I/O memory references	M/sec	1.946	0.936	0.481	0.458	6.260
I/O memory conflicts	Avg/ref	0.328	0.043	0.131	0.277	0.828

**Table 3: ACSF C90 Eagle 1H94 Daily Average HPM Measurements-
Global Counters**

Measurement	Unit	Avg	STD	COV	Min	Max
CPU time	Sec	517352.	106606.	0.206	177006.	652927.
Instruction Issue	M/sec	55.857	3.393	0.061	45.472	62.921
Average clock periods/inst	-----	4.313	0.268	0.062	3.814	5.278
CP holding issue	Percent	68.422	2.127	0.031	64.315	74.184
Instruction buffer fetches	M/sec	0.242	0.055	0.229	0.146	0.435
Floating Pt. Ops per CPU	M/sec	259.090	25.393	0.098	204.264	329.113
Vector Fl. Pt. Ops per CPU	M/sec	255.971	25.900	0.101	199.765	327.871
CPU memory references	M/sec	256.455	23.612	0.092	199.626	321.716
CPU memory conflicts	Avg/ref	0.410	0.035	0.086	0.340	0.506
VEC memory references	M/sec	251.142	24.304	0.097	193.568	317.786
B/T memory references	M/sec	1.334	0.310	0.233	0.643	2.180
I/O memory references	M/sec	5.151	1.974	0.383	1.601	10.579
I/O memory conflicts	Avg/ref	0.581	0.049	0.085	0.476	0.704

The large variation in the CPU time measurement reflects the requirement that the HPM data represent a continuous interval. Occasionally, persistent hardware and/or software problems may require several shutdowns during the 24-hour measurement period. The CPU time reported in the table for such days is the longest continuous period without a shutdown. The average VN CPU time is about twice as large as the average Eagle CPU time because VN has twice as many processors.

The tables show that the rate of instruction issue for the two machines corresponds to about 1 instruction every 4.5 clock periods. The time between instruction issues indicates an average period in which the operations produced by the instructions are being carried out and longer periods tend to characterize operations carried out by vector instructions. Pure scalar codes would issue about 1 instruction every clock period and highly vectorized CFD workload applications issue 1 instruction every 7 clock periods. The percent of hold issue CPs seems large, but the analysis following Table 11 will show that other operations were in progress during this time. The low rate of instruction buffer fetches indicates that the processors were busy executing code which generally kept the instruction buffer filled.

Both workloads are performing at modestly high FLOP rates and such performance indicate that the vector instructions are performing many operations. Although the average CPU rate is well below the single CPU maximum of 960 MFLOPS, many NAS applications display performance rates exceeding 500 MFLOPS. The ratio of vector FLOPS to total FLOPS is 0.987 for both workloads, so scalar FLOPS constitute about 1% of the total workload FLOPS.

The CPU MFLOP rate slightly exceeds the CPU memory reference rate, indicating that the average floating point operation must be reusing data in the registers to avoid memory accesses. The CPU memory reference rate is 18% of the 1440 M/sec memory bandwidth (6 words per CP). The memory conflict measurements indicate the average delay (in CPs) experienced by a typical memory access. A vector memory reference with no delay requires 1 CP to complete. For VN, this delay is 0.300 CP and for Eagle, this delay is 0.410.

The two workloads differ strongly in the CPU I/O memory reference rate, with VN displaying 1.95 Mwords/sec per CPU and Eagle showing 5.15 Mwords/sec per CPU. Monitoring of the VN I/O to the disks in 1H94 indicate an average transfer rate of 3.0 Mwords/sec and a maximum of 5.0 Mwords/sec for VN. For Eagle, the corresponding rates were 11.0 and 21.7 Mwords/sec. Since the HPM indicates a sustained I/O rate of 32MW/sec for VN and 41 MW/sec for Eagle, 90% of the C90 I/O targets the SSD while 66% of the Eagle I/O targets the SSD. Although the average I/O rate is well below the single CPU maximum of 239 Mwords/sec, several NAS applications sustain data transfer rates of 200 Mwords/sec. Typically, programs representing chemistry applications store and reuse a considerable amount of computed data and thus provide the highest I/O demand among the Ames C90 user community.

The I/O rate measurements display a large COV relative to the performance rate measurements. This variance reflects the differing input/output requirements of NAS users.

The Cray timesharing architecture decouples the I/O rate from the MFLOP rate because the data transfer occurs when the user program has given up control of the CPU to another program. The second program can maintain the CPU MFLOP rate while the I/O from the first program proceeds. If the transfer is efficient and the two programs have similar performance characteristics, measurements should show the MFLOP rate relatively constant while the I/O rate fluctuates according to user needs. The C90 measurements substantiate this claim.

3.0 Instruction Holds

Instructions are fetched from the instruction buffer by the instruction processor. If any of the resources required to execute the instruction are reserved, the instruction issue logic prevents the instruction from issuing. The HPM records all CPs for which the instruction holds issue and the table presents these as the percent of total CPU time. Since there may be more than one resource reservation preventing an instruction issue, the sum of the percentages in this group can exceed 100%.

**Table 4: NAS C90 VN 1H94 Daily Average HPM Measurements-
Instruction Holds**

Measurement	Unit	Avg	STD	COV	Min	Max
Waiting on A-registers	% CPU	4.841	0.343	0.071	3.990	5.602
Waiting on S-registers	% CPU	8.149	1.246	0.153	5.218	11.780
Waiting on V-registers	% CPU	23.892	1.610	0.067	19.364	27.333
Waiting on B/T-registers	% CPU	1.166	0.129	0.110	0.807	1.529
Waiting on F'nctnal Units	% CPU	26.134	1.887	0.072	21.474	31.798
Waiting on Shared Regs	% CPU	0.478	0.348	0.727	0.022	1.883
Waiting on Memory Ports	% CPU	17.686	2.064	0.117	13.235	23.464
Waiting on Miscellaneous	% CPU	2.410	0.114	0.047	2.098	2.724

**Table 5: ACSF C90 Eagle 1H94 Daily Average HPM Measurements-
Instruction Holds**

Measurement	Unit	Avg	STD	COV	Min	Max
Waiting on A-registers	% CPU	5.642	0.705	0.125	3.854	7.420
Waiting on S-registers	% CPU	8.783	1.607	0.183	4.972	14.208
Waiting on V-registers	% CPU	19.101	1.709	0.089	14.735	23.408
Waiting on B/T-registers	% CPU	1.223	0.184	0.150	0.806	1.750
Waiting on F'nctnal Units	% CPU	22.983	2.289	0.100	17.058	27.824
Waiting on Shared Regs	% CPU	0.024	0.035	1.453	0.001	0.214
Waiting on Memory Ports	% CPU	20.128	1.921	0.095	15.617	25.269
Waiting on Miscellaneous	% CPU	2.398	0.151	0.063	2.087	2.720

The major reason for instruction issue delays are busy vector registers and busy vector functional units. The instruction processor will not issue an instruction until operations in these units have completed. Calculations derived from counter data (Tables 10 and 11) have shown that other operations were in progress during these delays.

The approximately equal delays in vector registers and vector functional units indicates efficient register use and overlapping of vector functional units. The ratio of the sum of A (Address) and S (Scalar) register holds to V (Vector) register and Functional Unit holds is 0.259 for VN and 0.342 for Eagle. The higher Eagle ratio indicates a somewhat larger scalar workload component.

Sets of shared registers couple the C90 CPUs for efficient synchronization of parallel tasks. NAS provides strong incentives in the form of discounted CPU time for users employing parallel processing on VN whereas the ACSF provides no such discounts on Eagle. Jobs employing multiple processors consumed almost 40% of the VN CPU time in 1H94. The larger amount of Shared Register hold issue on VN reflects this usage. A handful of large projects executing in a special off-prime NQS queue account for most of this CPU time.

A CPU memory port accesses a section which accesses a memory bank. Memory references can lead to two kinds of delay in the C90 architecture. A memory instruction hold occurs, for example, when a register is reserved by another instruction or a memory port is busy. A memory conflict occurs when a needed bank is busy. A user program executing on a single CPU can encounter conflicts when it continuously references the same bank. A workload can encounter conflicts when several CPUs simultaneously reference the same bank.


The rate for memory transfer depends upon vector length because longer vector lengths (up to the hardware maximum of 128) can amortize the startup overhead. Tables 8 and 9 show that the average vector length is about 69 for both workloads. At this vector length, the C90 can store data at a rate of 1.28 CP/word. Data from Tables 2 and 3 indicate that each memory reference on the average experiences a memory contention delay of 0.300 CP for VN and 0.410 for Eagle. Table 4 and Table 5 data indicate that reserved memory resources prevent the CPU from issuing an instruction about 18% of the time for VN and 20% of the time for Eagle. Converting this aggregate delay to a delay per reference yields a instruction issue memory delay of 0.161 CP for VN and 0.192 CP for Eagle.

For VN, total memory delay is $0.300 + 0.161$, or about 0.461 CP/reference. For Eagle, total memory delay is $0.410 + 0.192$, or about 0.602 CP/reference. The average delay is a fraction of the 1.27 CP minimum required for the average workload vector memory reference.

The following figure shows 1994 C90 performance as a function of total memory delay and indicates a slight decrease in memory delay as VN performance increases and a constant delay as Eagle performance increases.

Memory Delay vs CPU Rate

NAS C90 1H94



The VN data indicate that reductions in memory delay accompany (or perhaps permit) increases in CPU performance, while the Eagle data indicates that memory delay does not decrease with increasing CPU performance. In practice, other factors such as the amount of vectorization and program vector length make the actual relationship between performance and its contributing factors multidimensional. HPM measurements indicate that other operations are in progress during these memory delays and these operations can offset the effect of the memory delays. The figure does indicate that memory delay does not increase as CPU performance increases and this result lends additional credibility to the observation that memory is not a bottleneck for these workloads.

4.0 Instruction Issues

Instructions produce the operations which constitute the actual workload tasks. The unit “M/sec” denotes “Millions of instructions per second”.

**Table 6: NAS C90 VN 1H94 Daily Average HPM Measurements-
Instruction Issues**

Measurement	Unit	Avg	STD	COV	Min	Max
(000-004)Special	M/sec	1.025	0.100	0.097	0.802	1.348
(005-017)Branch	M/sec	2.154	0.212	0.099	1.721	2.753
(02x,030-033)A Register	M/sec	24.249	1.618	0.067	19.730	32.412
(034-037)B/T Memory	M/sec	0.138	0.025	0.182	0.071	0.219
(040-043,071-077)S Register	M/sec	6.604	0.960	0.145	3.954	9.407
(044-061)Scalar Integer	M/sec	4.206	0.583	0.139	2.805	5.808
(062-070)Scalar Floating Pt.	M/sec	3.410	0.676	0.198	1.907	5.400
(10x-13x)Scalar Memory	M/sec	3.364	0.560	0.166	2.146	5.188
(140-177)All Vector	M/sec	8.433	0.475	0.056	7.282	10.045

**Table 7: ACSF C90 Eagle 1H94 Daily Average HPM Measurements-
Instruction Issues**

Measurement	Unit	Avg	STD	COV	Min	Max
(000-004)Special	M/sec	1.042	0.155	0.149	0.752	1.444
(005-017)Branch	M/sec	2.584	0.275	0.106	2.000	3.373
(02x,030-033)A Register	M/sec	25.669	1.511	0.059	22.995	31.089
(034-037)B/T Memory	M/sec	0.160	0.039	0.245	0.066	0.273
(040-043,071-077)S Register	M/sec	6.742	1.229	0.182	3.194	10.074
(044-061)Scalar Integer	M/sec	4.573	1.016	0.222	2.570	7.692
(062-070)Scalar Floating Pt.	M/sec	3.120	0.845	0.271	1.072	4.944
(10x-13x)Scalar Memory	M/sec	3.979	0.738	0.185	2.393	6.015
(140-177)All Vector	M/sec	7.990	0.601	0.075	6.271	9.395

For both workloads, A-register instructions comprise about 45% of the scalar instructions issued. These instructions compute memory addresses and indexes for memory, loop control, and I/O. All CPUs of the C90 architecture have two pipes, one consisting of an add functional unit and the other consisting of a multiply functional unit. The C90 functional units consist of 64 double-width functional units and this arrangement requires some additional A-register operations.

Scalar instructions constitute about 33% of workload instructions for both machines. Scalar instructions produce scalar operations. The scalar floating point rate, when combined with the vector floating point operation rate (Tables 2 and 3), gives the total floating point operation rate. The scalar floating point calculation are about 1% of total workload FLOPS.

Vector instructions are only 17% of the total instructions, but vector operations represent about 92% of the workload operations (Table 10 and 11). A single vector instruction can produce many vector operations and the term vector instruction denotes the average number of vector operations produces by a vector operation.

5.0 Vector Operations

All of the vector operations shown in Tables 8 and 9 are produced by vector instructions. Tables 6 and 7 show that the rate of instruction issue for all vector instructions was 8.433 million per second for VN and 7.990 for Eagle.

The vector operation rate for 1H94, which is the sum of the column 3 values in the first 8 rows of Tables 8 and 9, was 589 million per second for VN and 545 million per second for Eagle.

**Table 8: NAS C90 VN 1H94 Daily Average HPM Measurements-
Vector Operations**

Measurement	Unit	Avg	STD	COV	Min	Max
Vector Logical	M/sec	35.284	4.755	0.135	24.738	45.532
Vector Shift/Pop/LZ	M/sec	8.923	0.993	0.111	6.977	11.825
Vector Integer Add	M/sec	18.519	2.784	0.150	12.399	26.043
Vector Floating Pt. Multiply	M/sec	134.405	11.620	0.086	108.251	170.892
Vector Floating Pt. Add	M/sec	120.748	11.567	0.096	97.907	166.578
Vector Floating Reciprocal	M/sec	8.784	1.348	0.153	6.185	13.342
Vector Memory Read	M/sec	181.722	13.553	0.075	145.895	226.575
Vector Memory Write	M/sec	80.827	6.089	0.075	64.754	101.118
Average Vector Length	-----	69.969	4.868	0.070	57.370	87.190

**Table 9: C90 Eagle 1H94 Daily Average HPM Measurements-
Vector Operations**


Measurement	Unit	Avg	STD	COV	Min	Max
Vector Logical	M/sec	20.851	4.182	0.201	10.733	34.839
Vector Shift/Pop/LZ	M/sec	6.062	1.122	0.185	3.247	8.919
Vector Integer Add	M/sec	11.940	2.448	0.205	5.903	20.497
Vector Floating Pt. Multiply	M/sec	129.786	12.787	0.099	103.444	167.335
Vector Floating Pt. Add	M/sec	121.855	13.463	0.110	91.706	156.709
Vector Floating Reciprocal	M/sec	4.330	0.710	0.164	2.266	6.486
Vector Memory Read	M/sec	178.224	17.871	0.100	137.122	226.786
Vector Memory Write	M/sec	72.918	7.376	0.101	55.896	92.383
Average Vector Length	-----	68.480	5.785	0.084	53.360	83.130

Vector memory load (read) rates are twice as large as vector memory store (write) rates. A FLOP requires, on the average, one memory reference, but it is more likely to be a load than a store. The C90 architecture provides each CPU with two double-width memory paths for loading data from memory and one memory path for storage; the architecture reserves the fourth memory path for I/O and instruction buffer transfers. The C90 provides a maximum memory bandwidth of 6 references per CP per CPU. Since the maximum CPU computational rate is 4 floating point operations per CP, the Cray design attempts to ensure CPU-intensive codes will not experience memory-starvation.

The HPM measurements indicate that the current workloads require an average CPU memory bandwidth of 1.0 references per CP and a maximum memory bandwidth of 1.3 references per CP. Some individual NAS applications require as many as 2.6 references per CP per CPU to maintain their performance rate.

The tables show that VN performs a much higher rate of vector logical operations than Eagle. Some advanced algorithms producing high rates of logical operations can occur in codes containing unstructured or sparse matrix solvers as well as grid generation codes.

The ratio of total vector operations to total vector instructions is the workload average vector length. For both machines, the 1H94 value is about 69 whereas the C90 hardware vector length is 128. The C90 vector length reported by the hardware monitor is the program logical vector length modulo 128. The following figure shows the relationship between the hardware vector length and the program vector length.



A single program, with one constant length loop dominating FLOP performance, would display a vector length of 69 for loop logical lengths of either 69 or 138. Workload measurements indicate that C-90 average vector lengths have historically ranged between 50 and 80. Thus, average NAS workload vector lengths inhabit the first slope in the figure and the average value is definitely 69 as opposed to 138.

Many individual programs compose the workload and some of these codes displaying short vector lengths consume a considerable amount of CPU time. While the algorithmic properties of some NAS codes lead to short vector lengths, insufficient vector lengths remain the most visible performance problem. NAS has begun to educate its user community about the coding steps required for increasing vector length because Cray has failed to provide semantic guidelines for most effective use of its performance enhancement tool, FPP. For codes written in the proper format, FPP can recast the code to increase the vector length. Unfortunately, there is no written description of the proper format and most user codes derive little benefit from this tool. Perhaps NAS should provide a description of the proper FPP format in places accessible to NAS users.

6.0 Derived Data

The table lists several quantities obtained through calculations with the counter data.

Table 10: NAS C90 VN 1H94 Daily Average HPM Measurements-Derived Data

Measurement	Unit	Avg	STD	COV	Min	Max
System Availability	Percent	0.945	0.020	0.021	0.800	0.970
System MFLOPS	M/sec	4044.656	350.635	0.087	3316.650	5263.200
Vector Operation Fraction	Percent	92.831	0.908	0.010	89.890	95.130
Scalar Operation Fraction	Percent	7.169	0.908	0.127	4.870	10.110
Vector Operation Rate	M/sec	589.213	42.295	0.072	478.590	726.670
Scalar Operation Rate	M/sec	45.149	3.403	0.075	35.730	53.810
Total Operation Rate	M/sec	634.362	39.802	0.063	532.410	763.850
Instruction Issue Fraction	Percent	23.606	1.041	0.044	20.184	25.823
Hold Issue Fraction	Percent	67.729	1.531	0.023	64.286	72.684
Null Instruction Fraction	Percent	8.764	0.587	0.067	6.926	9.972

Table 11: ACSF C90 Eagle 1H94 Daily Average HPM Measurements-Derived Data

Measurement	Unit	Avg	STD	COV	Min	Max
System Availability	Percent	0.803	0.077	0.096	0.480	0.940
System MFLOPS	M/sec	1658.084	186.929	0.113	1095.640	2202.890
Vector Operation Fraction	Percent	91.847	1.204	0.013	89.000	94.780
Scalar Operation Fraction	Percent	8.153	1.204	0.148	5.220	11.000
Vector Operation Rate	M/sec	545.967	50.579	0.093	428.450	677.520
Scalar Operation Rate	M/sec	47.867	3.548	0.074	37.320	54.570
Total Operation Rate	M/sec	593.834	47.595	0.080	481.420	717.430
Instruction Issue Fraction	Percent	23.471	0.097	0.042	21.594	25.440
Hold Issue Fraction	Percent	66.514	1.599	0.024	63.360	69.770
Null Instruction Fraction	Percent	10.015	0.667	0.067	8.636	11.348

Availability is the fraction of time the C90 operated in user mode. During other times, the C90 was either idle or executing system calls. The next section will discuss the reasons for Eagle's much larger idle.

System MFLOPS denotes the system throughput. This rate is the product:

$$\text{System MFLOPS} = \text{MFLOPS/CPU} * \text{CPUs} * \text{Availability}.$$

The table shows the VN throughput rate to be 4045 MFLOPS (26.3% of peak) and the Eagle throughput to be 1658 MFLOPS (21.6% of peak). The major reason for the difference in performance of the two machines is the higher VN availability.

VN's slightly higher vector fraction produces a total operation rate which is about 6% greater than Eagle's. Both rates exceed 2 OPS/CP (VN's is 2.48 OP/CP), the instruction processor is able to overlap operations despite the large number of instruction hold issue CPs discussed under Table 1.

7.0 Discussion

The CPU memory conflicts (Table 2 and Table 3) and the availability (Table 10 and Table 11) were the major differences in the two workloads. This section discusses these differences and summarizes the performance histories of the two machines.

The HPM measurements showed that relative to VN, an average memory reference on Eagle experienced a 33% longer delay due to memory conflict. A subsequent calculation including instruction issue delay translated into a 20% longer memory delay for a typical memory reference. While neither one of these workloads is memory-bound, the high degree of similarity observed for the other workload measurements highlights this memory conflict discrepancy. The Eagle workload could contain some poorly written codes whose only signature in the workload is a large number of memory conflicts. It is more likely that, since the number of memory banks is a key factor in memory conflicts, and since Eagle has smaller number of memory banks (half the memory banks of VN), the hardware plays a role in the higher Eagle memory conflict ratio.

VN reported an 14% higher availability than Eagle. The following table gives the components of elapsed time during 1H94:

Table 12: NAS C90 1H94-Elapsed Time Components (Percentages)

Component	VN	Eagle
User	94.4	80.3
System	4.4	8.6
Idle	1.1	11.1

The table shows that VN had only 1% idle time while Eagle displayed 11% idle time. The VN workload arises from users throughout the country giving rise to a strong interactive component for about 12 hours each weekday. The Eagle workload is more local and the strong interactive component is present for about 8 hours. To offset potential idle, VN employs a set of deferred queues which have reduced charges and which are turned on only during times of low batch activity. VN also has a larger memory which makes it easier to service a workload with wide range of job memory requirements. Eagle has a smaller memory and no incentive for users to submit jobs during times of low batch activity. The Eagle idle is an administrative problem having a variety of possible remedies. Ideally, the remedy chosen to increase the Eagle availability should try to maintain VN availability and should impact all projects in a fair manner.

In 1H94, VN employed the UNICOS operating system version 8.0 while Eagle used version 7.C.2. Kernel contention, i.e., the updating of kernel data tables by a single CPU which prevents other CPUs from accessing those same tables is reduced by version 8.0. Kernel contention was one reason for the greater Eagle system time. The other reason was probably the Eagle's higher I/O rate.

The 1H94 average VN CPU performance was about 7% greater than that of 2H93. The following table summarizes some key results:

Table 13: NAS C90 Key Hardware Performance Results

Measurement	2Q93	3Q93	4Q93	1Q94	2Q94
CPU MFLOPS	244	244	255	274	261
Percent Vectorization	91.7	91.0	91.8	92.8	92.8
Vector Length	62.9	59.7	67.6	68.3	71.8
System Availability	85.6	88.3	88.7	94.8	94.3
System GFLOPS	3.315	3.442	3.626	4.161	3.933
System Efficiency	21.6	22.4	23.6	27.1	25.6

The table shows increased CPU MFLOPS and System GFLOPS during 1H94. At a 90% confidence level, the 1H94 confidence intervals for these two quantities lie outside these quantities' confidence intervals for the previous two quarters. Thus, the 1H94 performance increases are statistically significant and deserve an explanation. No hardware upgrade occurred during 1H94, but Cray upgraded the default Fortran compiler several times during this period. The following table presents the VN workload performance as a function of the default compiler.

Table 14: NAS VN Workload Performance History

Compiler Installation Date	Compiler Version	Days as System Default	Average Daily Mflops	Average Vector Length	Average Vector Fraction
04/29/93	5.0.4.13	17	271	63.0	91.5
05/18/93	5.0.4.17	58	252	57.2	91.0
07/14/93	6.0.0.0	56	243	59.3	91.1
09/09/93	6.0.0.4	67	241	63.9	87.4
11/15/93	6.0.0.9	76	262	67.2	92.1
02/04/94	6.0.2.3	105	268	69.9	93.0
05/20/94	6.0.3.5	41	261	72.8	92.9

This table shows a pronounced drop in CPU performance at the installation of Version 6.0.0.0 and a subsequent recovery in performance when Version 6.0.0.9 became the default

compiler. The data suggest that later releases of the 6.0 compiler corrected some of the initial optimization problems and these corrections led to the statistically significant performance increases observed for the workload

The Eagle workload data do not include the period in which early releases of the 6.0 compiler were the default, and measurements during 1H94 saw only version 6.0.2.3 of the 6.0 compiler.

8.0 Conclusion

The NAS C90 Von Neumann (VN) and the ACSF C90 Eagle computational workloads displayed similar CPU performance during 1H94. The composition of the user source code allowed the Cray compiler to generate machine code producing 91% vector operations. CPU FLOP rates averaged about 25% of peak.

Memory does not appear to be a bottleneck for these workloads. The ratio of floating point operations to memory operations was about 1.0. Analysis of the memory-related delays indicated no trends to increased memory delays as CPU performance increased. While Eagle displayed 25% more memory-induced delay than VN, its CPU performance was only 3% less than that of VN.

The dominant causes of instruction hold issues were the reservations placed on the vector units. Since the number of operations per clock period exceeded 1.0, other operations, such as calculations in the functional units were in progress during these periods of instruction hold.

Insufficient vector lengths continued to hinder the performance of both machines. While advanced algorithms such as multigrid relaxation schemes and unstructured grid solvers may display inherently short vector lengths, NAS should continue to encourage longer vector lengths for the other numerical techniques which constitute the bulk of the workload. To assist in the generation of longer vector lengths, NAS can promote intelligent use of the FPP tool, with a first step being the production of guidelines for writing Fortran code in form which FPP can optimize.

The ACSF CPUs experienced considerably more I/O traffic than VN, but the higher I/O rate did not correlate with decreased CPU performance.

The NAS C90 displayed a factor of 2.4 greater system throughput than the ACSF C90, This factor is somewhat larger than the factor of 2.0 expected on the basis of the number of CPUs. The smaller computational load experienced by the ACSF C90 led to a larger idle relative to the NAS C90 and NAS administrators may address this discrepancy in the future.

9.0 Acknowledgment

Thanks to David Barkai and Alan Powers for reviewing this paper.

10.0 References

- [1] Cray Research, Inc. UNICOS Performance Utilities Reference Manual, SR-2040 7.0, 1992.
- [2] Bergeron, R.J. "Performance Analysis of the NAS Eagle Workload", NASA Ames Research Center, NAS Systems Division Report RND-90-009, December, 1990.

The Performance of the NAS HSPs in 1st Half of 1994

Robert J. Bergeron

Report NAS-95-008 February 1995

NAS Systems Development Branch
NAS Systems Division
NASA Ames Research Center
Mail Stop 258-6
Moffett Field, CA 94035-1000

Abstract

During the first 6 months of 1994, the NAS 16-CPU Y-MP C90 Von Neumann (VN) delivered an average throughput of 4.045 GFLOPS while the ACSF 8-CPU Y-MP C90 Eagle averaged 1.658 GFLOPS. The VN rate represents a machine efficiency of 26.3% whereas the Eagle rate corresponds to a machine efficiency of 21.6%. VN displayed a greater efficiency than Eagle primarily because the stronger workload demand for its CPU cycles allowed it to devote more time to user programs and less time to idle. An additional factor increasing VN efficiency was the ability of the UNICOS 8.0 Operating System to deliver a larger fraction of CPU time to user programs. Although measurements indicate increasing vector length for both workloads, insufficient vector lengths continue to hinder HSP performance. To improve HSP performance, NAS should continue to encourage the HSP users to modify their codes to increase program vector length.

1.0 Introduction

The introduction of the C90 in March 1993 motivated the daily monitoring of the hardware performance of the NAS High Speed Processors (HSPs). The C90 Hardware Performance Monitor (HPM) continuously delivers a full 32-counter record for the workload [1]. NAS records the daily average values of all HPM counters and this paper, covering the 1st half of 1994, is the fourth in a series of reports using these counters to evaluate the performance of the NAS workload.

A NASA Ames administrative action brought the Aeronautics Consolidated Supercomputer Facility (ACSF) C90 Eagle under control of NAS in the second half of 1993. Daily HPM monitoring of the ACSF C90 Eagle began in late March of 1994. The following table presents the characteristics of the two machines.

Table 1: Characteristics of NAS HSPs

Characteristic	Unit	Von Neumann	Eagle
Serial Number	---	4012	4015
Number of CPUs	---	16	8
Clock Cycle	ns	4.167	4.167
Memory	MW	1024	128
Memory Banks	---	1024	512
Memory Bank Cycle Time	Clock Period	23	23
SSD Size	MW	1024	512
VHISPs	---	4	2
IOS	---	Model E	Model E
UNICOS	Version	8.0.1	7.C.2

This report provides tables of counter values representing the average, maximum, and minimum values from the daily reports in the first half of 1994. The NAS C90 Von Neumann(VN) provided 180 such daily reports

whereas the ACSF Eagle provided 105 daily reports since monitoring began in the middle of the first half of 1994.

The counter value tables provide performance rate data per CPU for the actual time the CPU spent executing the user programs. System throughput is derived from the CPU Floating Point Operation (FLOP) rate, the wall clock time, and the total number of CPUs. A complete explanation of all counter data occurs in [2].

To provide a feel for the daily variation in each of the counters, the report also provides the standard deviation (STD) and coefficient of variation (COV). The coefficient of variation is the ratio of the standard deviation of a quantity divided by its average value.

The report divides the 32 C90 counters into 4 functional groups: global counters, instruction holds, instruction issues, and vector operations. Sections 2 through 5 describe the results obtained from each of the four groups and compares the measurements from the two workloads.

2.0 Global Counter Data

Table 2 provides counter data giving a total counts for instructions, operations and references. The unit "M/sec" denotes "Million per sec" and the unit "avg/ref" denotes "average (conflict) per reference". The term "reference" denotes a single Cray word (8-byte) data transfer.

Table 2: NAS C90 VN 1H94 Daily Average HPM Measurements-Global Counters

Measurement	Unit	Avg	STD	COV	Min	Max
CPU time	Sec	1096107.	277109.	0.253	155327.	1531277.
Instruction Issue	M/sec	53.583	3.379	0.063	43.297	61.972
Average clock periods/inst	-----	4.497	0.287	0.064	3.872	5.543
CP holding issue	Percent	69.716	2.019	0.029	63.776	75.239
Instruction buffer fetches	M/sec	0.273	0.039	0.141	0.178	0.385
Floating Pt. Ops per CPU	M/sec	267.347	22.475	0.084	218.794	346.303
Vector Fl. Pt. Ops per CPU	M/sec	263.937	22.838	0.087	213.743	344.396
CPU memory references	M/sec	267.175	18.371	0.069	217.479	330.843
CPU memory conflicts	Avg/ref	0.300	0.049	0.164	0.226	0.643
VEC memory references	M/sec	262.550	18.936	0.072	210.648	327.693
B/T memory references	M/sec	1.262	0.219	0.173	0.712	1.974
I/O memory references	M/sec	1.946	0.936	0.481	0.458	6.260
I/O memory conflicts	Avg/ref	0.328	0.043	0.131	0.277	0.828

Table 3: ACSF C90 Eagle 1H94 Daily Average HPM Measurements-Global Counters

Measurement	Unit	Avg	STD	COV	Min	Max
CPU time	Sec	517352.	106606.	0.206	177006.	652927.
Instruction Issue	M/sec	55.857	3.393	0.061	45.472	62.921
Average clock periods/inst	-----	4.313	0.268	0.062	3.814	5.278
CP holding issue	Percent	68.422	2.127	0.031	64.315	74.184
Instruction buffer fetches	M/sec	0.242	0.055	0.229	0.146	0.435
Floating Pt. Ops per CPU	M/sec	259.090	25.393	0.098	204.264	329.113
Vector Fl. Pt. Ops per CPU	M/sec	255.971	25.900	0.101	199.765	327.871
CPU memory references	M/sec	256.455	23.612	0.092	199.626	321.716
CPU memory conflicts	Avg/ref	0.410	0.035	0.086	0.340	0.506
VEC memory references	M/sec	251.142	24.304	0.097	193.568	317.786
B/T memory references	M/sec	1.334	0.310	0.233	0.643	2.180
I/O memory references	M/sec	5.151	1.974	0.383	1.601	10.579
I/O memory conflicts	Avg/ref	0.581	0.049	0.085	0.476	0.704

The large variation in the CPU time measurement reflects the requirement that the HPM data represent a continuous interval. Occasionally, persistent hardware and/or software problems may require several shutdowns during the 24-hour measurement period. The CPU time reported in the table for such days is the longest continuous period without a shutdown. The average VN CPU time is about twice as large as the average Eagle CPU time because VN has twice as many processors.

The tables show that the rate of instruction issue for the two machines corresponds to about 1 instruction every 4.5 clock periods. The time between instruction issues indicates an average period in which the operations produced by the instructions are being carried out and longer periods tend to characterize operations carried out by vector instructions. Pure scalar codes would issue about 1 instruction every clock period and highly vectorized CFD workload applications issue 1 instruction every 7 clock periods. The percent of hold issue CPs seems large, but the analysis following Table 11 will show that other operations were in progress during this time. The low rate of instruction buffer fetches indicates that the processors were busy executing code which generally kept the instruction buffer filled.

Both workloads are performing at modestly high FLOP rates and such performance indicate that the vector instructions are performing many operations. Although the average CPU rate is well below the single CPU maximum of 960 MFLOPS, many NAS applications display performance rates exceeding 500 MFLOPS. The ratio of vector FLOPS to total FLOPS is 0.987 for both workloads, so scalar FLOPS constitute about 1% of the total workload FLOPS.

The CPU MFLOP rate slightly exceeds the CPU memory reference rate, indicating that the average floating point operation must be reusing data in the registers to avoid memory accesses. The CPU memory reference rate is 18% of the 1440 M/sec memory bandwidth (6 words per CP). The memory conflict measurements indicate the average delay (in CPs) experienced by a typical memory access. A vector memory reference with no delay requires 1 CP to complete. For VN, this delay is 0.300 CP and for Eagle, this delay is 0.410.

The two workloads differ strongly in the CPU I/O memory reference rate, with VN displaying 1.95 Mwords/sec per CPU and Eagle showing 5.15 Mwords/sec per CPU. Monitoring of the VN I/O to the disks in 1H94 indicate an average transfer rate of 3.0 Mwords/sec and a maximum of 5.0 Mwords/sec for VN. For Eagle, the corresponding rates were 11.0 and 21.7 Mwords/sec. Since the HPM indicates a sustained I/O rate of 32MW/sec for VN and 41 MW/sec for Eagle, 90% of the C90 I/O targets the SSD while 66% of the Eagle I/O targets the SSD. Although the average I/O rate is well below the single CPU maximum of 239 Mwords/sec, several NAS applications sustain data transfer rates of 200 Mwords/sec. Typically, programs representing chemistry applications store and reuse a considerable amount of computed data and thus provide the highest I/O demand among the Ames C90 user community.

The I/O rate measurements display a large COV relative to the performance rate measurements. This variance reflects the differing input/output requirements of NAS users.

The Cray timesharing architecture decouples the I/O rate from the MFLOP rate because the data transfer occurs when the user program has given up control of the CPU to another program. The second program can maintain the CPU MFLOP rate while the I/O from the first program proceeds. If the transfer is efficient and the two programs have similar performance characteristics, measurements should show the MFLOP rate relatively constant while the I/O rate fluctuates according to user needs. The C90 measurements substantiate this claim.

3.0 Instruction Holds

Instructions are fetched from the instruction buffer by the instruction processor. If any of the resources required to execute the instruction are reserved, the instruction issue logic prevents the instruction from issuing. The HPM records all CPs for which the instruction holds issue and the table presents these as the percent of total CPU time. Since there may be more than one resource reservation preventing an instruction issue, the sum of the percentages in this group can exceed 100%.

**Table 4: NAS C90 VN 1H94 Daily Average HPM Measurements-
Instruction Holds**

Measurement	Unit	Avg	STD	COV	Min	Max
Waiting on A-registers	% CPU	4.841	0.343	0.071	3.990	5.602
Waiting on S-registers	% CPU	8.149	1.246	0.153	5.218	11.780
Waiting on V-registers	% CPU	23.892	1.610	0.067	19.364	27.333
Waiting on B/T-registers	% CPU	1.166	0.129	0.110	0.807	1.529
Waiting on F'nctnal Units	% CPU	26.134	1.887	0.072	21.474	31.798
Waiting on Shared Regs	% CPU	0.478	0.348	0.727	0.022	1.883
Waiting on Memory Ports	% CPU	17.686	2.064	0.117	13.235	23.464
Waiting on Miscellaneous	% CPU	2.410	0.114	0.047	2.098	2.724

**Table 5: ACSF C90 Eagle 1H94 Daily Average HPM Measurements-
Instruction Holds**

Measurement	Unit	Avg	STD	COV	Min	Max
Waiting on A-registers	% CPU	5.642	0.705	0.125	3.854	7.420
Waiting on S-registers	% CPU	8.783	1.607	0.183	4.972	14.208
Waiting on V-registers	% CPU	19.101	1.709	0.089	14.735	23.408
Waiting on B/T-registers	% CPU	1.223	0.184	0.150	0.806	1.750
Waiting on F'nctnal Units	% CPU	22.983	2.289	0.100	17.058	27.824
Waiting on Shared Regs	% CPU	0.024	0.035	1.453	0.001	0.214
Waiting on Memory Ports	% CPU	20.128	1.921	0.095	15.617	25.269
Waiting on Miscellaneous	% CPU	2.398	0.151	0.063	2.087	2.720

The major reason for instruction issue delays are busy vector registers and busy vector functional units. The instruction processor will not issue an instruction until operations in these units have completed. Calculations derived from counter data (Tables 10 and 11) have shown that other operations were in progress during these delays.

The approximately equal delays in vector registers and vector functional units indicates efficient register use and overlapping of vector functional units. The ratio of the sum of A (Address) and S (Scalar) register holds to V (Vector) register and Functional Unit holds is 0.259 for VN and 0.342 for Eagle. The higher Eagle ratio indicates a somewhat larger scalar workload component.

Sets of shared registers couple the C90 CPUs for efficient synchronization of parallel tasks. NAS provides strong incentives in the form of discounted CPU time for users employing parallel processing on VN whereas the ACSF provides no such discounts on Eagle. Jobs employing multiple processors consumed almost 40% of the VN CPU time in 1H94. The larger amount of Shared Register hold issue on VN reflects this usage. A handful of large projects executing in a special off-prime NQS queue account for most of this CPU time.

A CPU memory port accesses a section which accesses a memory bank. Memory references can lead to two kinds of delay in the C90 architecture. A memory instruction hold occurs, for example, when a register is reserved by another instruction or a memory port is busy. A memory conflict occurs when a needed bank is busy. A user program executing on a single CPU can encounter conflicts when it continuously references the same bank. A workload can encounter conflicts when several CPUs simultaneously reference the same bank.


The rate for memory transfer depends upon vector length because longer vector lengths (up to the hardware maximum of 128) can amortize the startup overhead. Tables 8 and 9 show that the average vector length is about 69 for both workloads. At this vector length, the C90 can store data at a rate of 1.28 CP/word. Data from Tables 2 and 3 indicate that each memory reference on the average experiences a memory contention delay of 0.300 CP for VN and 0.410 for Eagle. Table 4 and Table 5 data indicate that reserved memory resources prevent the CPU from issuing an instruction about 18% of the time for VN and 20% of the time for Eagle. Converting this aggregate delay to a delay per reference yields a instruction issue memory delay of 0.161 CP for VN and 0.192 CP for Eagle.

For VN, total memory delay is $0.300 + 0.161$, or about 0.461 CP/reference. For Eagle, total memory delay is $0.410 + 0.192$, or about 0.602 CP/reference. The average delay is a fraction of the 1.27 CP minimum required for the average workload vector memory reference.

The following figure shows 1994 C90 performance as a function of total memory delay and indicates a slight decrease in memory delay as VN performance increases and a constant delay as Eagle performance increases.

Memory Delay vs CPU Rate

NAS C90 1H94



The VN data indicate that reductions in memory delay accompany (or perhaps permit) increases in CPU performance, while the Eagle data indicates that memory delay does not decrease with increasing CPU performance. In practice, other factors such as the amount of vectorization and program vector length make the actual relationship between performance and its contributing factors multidimensional. HPM measurements indicate that other operations are in progress during these memory delays and these operations can offset the effect of the memory delays. The figure does indicate that memory delay does not increase as CPU performance increases and this result lends additional credibility to the observation that memory is not a bottleneck for these workloads.

4.0 Instruction Issues

Instructions produce the operations which constitute the actual workload tasks. The unit “M/sec” denotes “Millions of instructions per second”.

**Table 6: NAS C90 VN 1H94 Daily Average HPM Measurements-
Instruction Issues**

Measurement	Unit	Avg	STD	COV	Min	Max
(000-004)Special	M/sec	1.025	0.100	0.097	0.802	1.348
(005-017)Branch	M/sec	2.154	0.212	0.099	1.721	2.753
(02x,030-033)A Register	M/sec	24.249	1.618	0.067	19.730	32.412
(034-037)B/T Memory	M/sec	0.138	0.025	0.182	0.071	0.219
(040-043,071-077)S Register	M/sec	6.604	0.960	0.145	3.954	9.407
(044-061)Scalar Integer	M/sec	4.206	0.583	0.139	2.805	5.808
(062-070)Scalar Floating Pt.	M/sec	3.410	0.676	0.198	1.907	5.400
(10x-13x)Scalar Memory	M/sec	3.364	0.560	0.166	2.146	5.188
(140-177)All Vector	M/sec	8.433	0.475	0.056	7.282	10.045

**Table 7: ACSF C90 Eagle 1H94 Daily Average HPM Measurements-
Instruction Issues**

Measurement	Unit	Avg	STD	COV	Min	Max
(000-004)Special	M/sec	1.042	0.155	0.149	0.752	1.444
(005-017)Branch	M/sec	2.584	0.275	0.106	2.000	3.373
(02x,030-033)A Register	M/sec	25.669	1.511	0.059	22.995	31.089
(034-037)B/T Memory	M/sec	0.160	0.039	0.245	0.066	0.273
(040-043,071-077)S Register	M/sec	6.742	1.229	0.182	3.194	10.074
(044-061)Scalar Integer	M/sec	4.573	1.016	0.222	2.570	7.692
(062-070)Scalar Floating Pt.	M/sec	3.120	0.845	0.271	1.072	4.944
(10x-13x)Scalar Memory	M/sec	3.979	0.738	0.185	2.393	6.015
(140-177)All Vector	M/sec	7.990	0.601	0.075	6.271	9.395

For both workloads, A-register instructions comprise about 45% of the scalar instructions issued. These instructions compute memory addresses and indexes for memory, loop control, and I/O. All CPUs of the C90 architecture have two pipes, one consisting of an add functional unit and the other consisting of a multiply functional unit. The C90 functional units consist of 64 double-width functional units and this arrangement requires some additional A-register operations.

Scalar instructions constitute about 33% of workload instructions for both machines. Scalar instructions produce scalar operations. The scalar floating point rate, when combined with the vector floating point operation rate (Tables 2 and 3), gives the total floating point operation rate. The scalar floating point calculation are about 1% of total workload FLOPS.

Vector instructions are only 17% of the total instructions, but vector operations represent about 92% of the workload operations (Table 10 and 11). A single vector instruction can produce many vector operations and the term vector instruction denotes the average number of vector operations produces by a vector operation.

5.0 Vector Operations

All of the vector operations shown in Tables 8 and 9 are produced by vector instructions. Tables 6 and 7 show that the rate of instruction issue for all vector instructions was 8.433 million per second for VN and 7.990 for Eagle.

The vector operation rate for 1H94, which is the sum of the column 3 values in the first 8 rows of Tables 8 and 9, was 589 million per second for VN and 545 million per second for Eagle.

**Table 8: NAS C90 VN 1H94 Daily Average HPM Measurements-
Vector Operations**

Measurement	Unit	Avg	STD	COV	Min	Max
Vector Logical	M/sec	35.284	4.755	0.135	24.738	45.532
Vector Shift/Pop/LZ	M/sec	8.923	0.993	0.111	6.977	11.825
Vector Integer Add	M/sec	18.519	2.784	0.150	12.399	26.043
Vector Floating Pt. Multiply	M/sec	134.405	11.620	0.086	108.251	170.892
Vector Floating Pt. Add	M/sec	120.748	11.567	0.096	97.907	166.578
Vector Floating Reciprocal	M/sec	8.784	1.348	0.153	6.185	13.342
Vector Memory Read	M/sec	181.722	13.553	0.075	145.895	226.575
Vector Memory Write	M/sec	80.827	6.089	0.075	64.754	101.118
Average Vector Length	-----	69.969	4.868	0.070	57.370	87.190

**Table 9: C90 Eagle 1H94 Daily Average HPM Measurements-
Vector Operations**


Measurement	Unit	Avg	STD	COV	Min	Max
Vector Logical	M/sec	20.851	4.182	0.201	10.733	34.839
Vector Shift/Pop/LZ	M/sec	6.062	1.122	0.185	3.247	8.919
Vector Integer Add	M/sec	11.940	2.448	0.205	5.903	20.497
Vector Floating Pt. Multiply	M/sec	129.786	12.787	0.099	103.444	167.335
Vector Floating Pt. Add	M/sec	121.855	13.463	0.110	91.706	156.709
Vector Floating Reciprocal	M/sec	4.330	0.710	0.164	2.266	6.486
Vector Memory Read	M/sec	178.224	17.871	0.100	137.122	226.786
Vector Memory Write	M/sec	72.918	7.376	0.101	55.896	92.383
Average Vector Length	-----	68.480	5.785	0.084	53.360	83.130

Vector memory load (read) rates are twice as large as vector memory store (write) rates. A FLOP requires, on the average, one memory reference, but it is more likely to be a load than a store. The C90 architecture provides each CPU with two double-width memory paths for loading data from memory and one memory path for storage; the architecture reserves the fourth memory path for I/O and instruction buffer transfers. The C90 provides a maximum memory bandwidth of 6 references per CP per CPU. Since the maximum CPU computational rate is 4 floating point operations per CP, the Cray design attempts to ensure CPU-intensive codes will not experience memory-starvation.

The HPM measurements indicate that the current workloads require an average CPU memory bandwidth of 1.0 references per CP and a maximum memory bandwidth of 1.3 references per CP. Some individual NAS applications require as many as 2.6 references per CP per CPU to maintain their performance rate.

The tables show that VN performs a much higher rate of vector logical operations than Eagle. Some advanced algorithms producing high rates of logical operations can occur in codes containing unstructured or sparse matrix solvers as well as grid generation codes.

The ratio of total vector operations to total vector instructions is the workload average vector length. For both machines, the 1H94 value is about 69 whereas the C90 hardware vector length is 128. The C90 vector length reported by the hardware monitor is the program logical vector length modulo 128. The following figure shows the relationship between the hardware vector length and the program vector length.



A single program, with one constant length loop dominating FLOP performance, would display a vector length of 69 for loop logical lengths of either 69 or 138. Workload measurements indicate that C-90 average vector lengths have historically ranged between 50 and 80. Thus, average NAS workload vector lengths inhabit the first slope in the figure and the average value is definitely 69 as opposed to 138.

Many individual programs compose the workload and some of these codes displaying short vector lengths consume a considerable amount of CPU time. While the algorithmic properties of some NAS codes lead to short vector lengths, insufficient vector lengths remain the most visible performance problem. NAS has begun to educate its user community about the coding steps required for increasing vector length because Cray has failed to provide semantic guidelines for most effective use of its performance enhancement tool, FPP. For codes written in the proper format, FPP can recast the code to increase the vector length. Unfortunately, there is no written description of the proper format and most user codes derive little benefit from this tool. Perhaps NAS should provide a description of the proper FPP format in places accessible to NAS users.

6.0 Derived Data

The table lists several quantities obtained through calculations with the counter data.

Table 10: NAS C90 VN 1H94 Daily Average HPM Measurements-Derived Data

Measurement	Unit	Avg	STD	COV	Min	Max
System Availability	Percent	0.945	0.020	0.021	0.800	0.970
System MFLOPS	M/sec	4044.656	350.635	0.087	3316.650	5263.200
Vector Operation Fraction	Percent	92.831	0.908	0.010	89.890	95.130
Scalar Operation Fraction	Percent	7.169	0.908	0.127	4.870	10.110
Vector Operation Rate	M/sec	589.213	42.295	0.072	478.590	726.670
Scalar Operation Rate	M/sec	45.149	3.403	0.075	35.730	53.810
Total Operation Rate	M/sec	634.362	39.802	0.063	532.410	763.850
Instruction Issue Fraction	Percent	23.606	1.041	0.044	20.184	25.823
Hold Issue Fraction	Percent	67.729	1.531	0.023	64.286	72.684
Null Instruction Fraction	Percent	8.764	0.587	0.067	6.926	9.972

Table 11: ACSF C90 Eagle 1H94 Daily Average HPM Measurements-Derived Data

Measurement	Unit	Avg	STD	COV	Min	Max
System Availability	Percent	0.803	0.077	0.096	0.480	0.940
System MFLOPS	M/sec	1658.084	186.929	0.113	1095.640	2202.890
Vector Operation Fraction	Percent	91.847	1.204	0.013	89.000	94.780
Scalar Operation Fraction	Percent	8.153	1.204	0.148	5.220	11.000
Vector Operation Rate	M/sec	545.967	50.579	0.093	428.450	677.520
Scalar Operation Rate	M/sec	47.867	3.548	0.074	37.320	54.570
Total Operation Rate	M/sec	593.834	47.595	0.080	481.420	717.430
Instruction Issue Fraction	Percent	23.471	0.097	0.042	21.594	25.440
Hold Issue Fraction	Percent	66.514	1.599	0.024	63.360	69.770
Null Instruction Fraction	Percent	10.015	0.667	0.067	8.636	11.348

Availability is the fraction of time the C90 operated in user mode. During other times, the C90 was either idle or executing system calls. The next section will discuss the reasons for Eagle's much larger idle.

System MFLOPS denotes the system throughput. This rate is the product:

$$\text{System MFLOPS} = \text{MFLOPS/CPU} * \text{CPUs} * \text{Availability}.$$

The table shows the VN throughput rate to be 4045 MFLOPS (26.3% of peak) and the Eagle throughput to be 1658 MFLOPS (21.6% of peak). The major reason for the difference in performance of the two machines is the higher VN availability.

VN's slightly higher vector fraction produces a total operation rate which is about 6% greater than Eagle's. Both rates exceed 2 OPS/CP (VN's is 2.48 OP/CP), the instruction processor is able to overlap operations despite the large number of instruction hold issue CPs discussed under Table 1.

7.0 Discussion

The CPU memory conflicts (Table 2 and Table 3) and the availability (Table 10 and Table 11) were the major differences in the two workloads. This section discusses these differences and summarizes the performance histories of the two machines.

The HPM measurements showed that relative to VN, an average memory reference on Eagle experienced a 33% longer delay due to memory conflict. A subsequent calculation including instruction issue delay translated into a 20% longer memory delay for a typical memory reference. While neither one of these workloads is memory-bound, the high degree of similarity observed for the other workload measurements highlights this memory conflict discrepancy. The Eagle workload could contain some poorly written codes whose only signature in the workload is a large number of memory conflicts. It is more likely that, since the number of memory banks is a key factor in memory conflicts, and since Eagle has smaller number of memory banks (half the memory banks of VN), the hardware plays a role in the higher Eagle memory conflict ratio.

VN reported an 14% higher availability than Eagle. The following table gives the components of elapsed time during 1H94:

Table 12: NAS C90 1H94-Elapsed Time Components (Percentages)

Component	VN	Eagle
User	94.4	80.3
System	4.4	8.6
Idle	1.1	11.1

The table shows that VN had only 1% idle time while Eagle displayed 11% idle time. The VN workload arises from users throughout the country giving rise to a strong interactive component for about 12 hours each weekday. The Eagle workload is more local and the strong interactive component is present for about 8 hours. To offset potential idle, VN employs a set of deferred queues which have reduced charges and which are turned on only during times of low batch activity. VN also has a larger memory which makes it easier to service a workload with wide range of job memory requirements. Eagle has a smaller memory and no incentive for users to submit jobs during times of low batch activity. The Eagle idle is an administrative problem having a variety of possible remedies. Ideally, the remedy chosen to increase the Eagle availability should try to maintain VN availability and should impact all projects in a fair manner.

In 1H94, VN employed the UNICOS operating system version 8.0 while Eagle used version 7.C.2. Kernel contention, i.e., the updating of kernel data tables by a single CPU which prevents other CPUs from accessing those same tables is reduced by version 8.0. Kernel contention was one reason for the greater Eagle system time. The other reason was probably the Eagle's higher I/O rate.

The 1H94 average VN CPU performance was about 7% greater than that of 2H93. The following table summarizes some key results:

Table 13: NAS C90 Key Hardware Performance Results

Measurement	2Q93	3Q93	4Q93	1Q94	2Q94
CPU MFLOPS	244	244	255	274	261
Percent Vectorization	91.7	91.0	91.8	92.8	92.8
Vector Length	62.9	59.7	67.6	68.3	71.8
System Availability	85.6	88.3	88.7	94.8	94.3
System GFLOPS	3.315	3.442	3.626	4.161	3.933
System Efficiency	21.6	22.4	23.6	27.1	25.6

The table shows increased CPU MFLOPS and System GFLOPS during 1H94. At a 90% confidence level, the 1H94 confidence intervals for these two quantities lie outside these quantities' confidence intervals for the previous two quarters. Thus, the 1H94 performance increases are statistically significant and deserve an explanation. No hardware upgrade occurred during 1H94, but Cray upgraded the default Fortran compiler several times during this period. The following table presents the VN workload performance as a function of the default compiler.

Table 14: NAS VN Workload Performance History

Compiler Installation Date	Compiler Version	Days as System Default	Average Daily Mflops	Average Vector Length	Average Vector Fraction
04/29/93	5.0.4.13	17	271	63.0	91.5
05/18/93	5.0.4.17	58	252	57.2	91.0
07/14/93	6.0.0.0	56	243	59.3	91.1
09/09/93	6.0.0.4	67	241	63.9	87.4
11/15/93	6.0.0.9	76	262	67.2	92.1
02/04/94	6.0.2.3	105	268	69.9	93.0
05/20/94	6.0.3.5	41	261	72.8	92.9

This table shows a pronounced drop in CPU performance at the installation of Version 6.0.0.0 and a subsequent recovery in performance when Version 6.0.0.9 became the default

compiler. The data suggest that later releases of the 6.0 compiler corrected some of the initial optimization problems and these corrections led to the statistically significant performance increases observed for the workload

The Eagle workload data do not include the period in which early releases of the 6.0 compiler were the default, and measurements during 1H94 saw only version 6.0.2.3 of the 6.0 compiler.

8.0 Conclusion

The NAS C90 Von Neumann (VN) and the ACSF C90 Eagle computational workloads displayed similar CPU performance during 1H94. The composition of the user source code allowed the Cray compiler to generate machine code producing 91% vector operations. CPU FLOP rates averaged about 25% of peak.

Memory does not appear to be a bottleneck for these workloads. The ratio of floating point operations to memory operations was about 1.0. Analysis of the memory-related delays indicated no trends to increased memory delays as CPU performance increased. While Eagle displayed 25% more memory-induced delay than VN, its CPU performance was only 3% less than that of VN.

The dominant causes of instruction hold issues were the reservations placed on the vector units. Since the number of operations per clock period exceeded 1.0, other operations, such as calculations in the functional units were in progress during these periods of instruction hold.

Insufficient vector lengths continued to hinder the performance of both machines. While advanced algorithms such as multigrid relaxation schemes and unstructured grid solvers may display inherently short vector lengths, NAS should continue to encourage longer vector lengths for the other numerical techniques which constitute the bulk of the workload. To assist in the generation of longer vector lengths, NAS can promote intelligent use of the FPP tool, with a first step being the production of guidelines for writing Fortran code in form which FPP can optimize.

The ACSF CPUs experienced considerably more I/O traffic than VN, but the higher I/O rate did not correlate with decreased CPU performance.

The NAS C90 displayed a factor of 2.4 greater system throughput than the ACSF C90, This factor is somewhat larger than the factor of 2.0 expected on the basis of the number of CPUs. The smaller computational load experienced by the ACSF C90 led to a larger idle relative to the NAS C90 and NAS administrators may address this discrepancy in the future.

9.0 Acknowledgment

Thanks to David Barkai and Alan Powers for reviewing this paper.

10.0 References

- [1] Cray Research, Inc. UNICOS Performance Utilities Reference Manual, SR-2040 7.0, 1992.
- [2] Bergeron, R.J. "Performance Analysis of the NAS Eagle Workload", NASA Ames Research Center, NAS Systems Division Report RND-90-009, December, 1990.